Searching for the Sources of Productivity

From macro to micro and back

Eric J. Bartelsman[1]

Vrije Universiteit Amsterdam; Tinbergen Institute; IZA Bonn

PRELIMINARY DRAFT

October 16, 2008

Abstract

Research into the sources of productivity growth has taken many disparate routes, from theoretical modeling using an aggregate perspective, to empirical explorations using firm-level data. This paper shows that models relying on the construct of the representative firm do not provide much scope to consider the effects of policy on productivity. The paper provides examples from the literature of the evidence against the representative firm assumption. Once the model is extended to include entry, exit, and heterogeneous firms, policy is able to affect aggregate productivity through many paths, including the efficiency of resource allocation and the margin of selection. As an example, some empirical explorations are provided for the effects of idiosyncratic distortions to firm profitability, the effects of exit costs, and the types of policy that may boost intangible investments.

[1] I am very indebted to my co-authors from the research projects cited in this paper: Pieter Gautier, John Halitwanger, Enrico Perotti, Stefano Scarpetta, and Joris de Wind, and all participants of the ONS/Eurostat ICT Impacts project.

1. Introduction

Research into the sources of productivity growth has taken many disparate routes, from theoretical modeling using an aggregate, or macro, perspective to empirical explorations using firm-level, or micro, data. Each route uses its own methodology to answer its own set of questions. Different approaches to research on a common topic, say the effect of R&D on growth, generally lead to answers that are not directly comparable as a consequence of the different strategies used to specify the problem and identify the results.

For policy makers, the state of affairs of productivity research is frustrating, at best. The main question on the table seems fairly straightforward: Which policies can be used to boost productivity (growth)? Yet, a review of the literature on productivity shows that available research does not allow one to prepare such a list of policy priorities. More specific questions, such as: "What is the impact of ICT on productivity" would seem easier to answer, but even here the answers depend greatly on the research methods and on the unit of analysis. And, even a specific result stating how much aggregate productivity would increase with a given increase in ICT use by firms provides little guidance on what type of policy change would have an effect on both ICT use and productivity.

This paper does not attempt to present a list of policy priorities in support of productivity growth. Instead, the paper will provide an annotated review of productivity research, discussing the mismatch in unit of analysis between theoretical and empirical work. Next, some recent empirical work using firm-level data is highlighted. A discussion then is provided of recent general equilibrium models of heterogeneous firms that track the effects of policy on individual firm choices as well as on the interplay between firm actions and market selection. Finally, some examples are given of recent empirical work that utilize both firm-level and aggregated data to study the impact of policy changes on intangible investment and aggregate productivity.

2. Theory and Empirics; Micro and Macro

This section provides some background on the theory underpinning much of the empirical productivity research. We use an example from recent productivity research, namely the impact of ICT on productivity, to show how and why traditional theory does not help in disentangling the sources of productivity growth. The main intuition is that while the macroeconomic growth theories generally appeal to microeconomic behavior of representative firms, the macro economy is made up of heterogeneous firms interacting in a market. A policy change that may have no theoretical link to productivity of a representative firm, may affect aggregate productivity through strategic interactions between individual firms. Alternatively, a policy expected to affect the behavior of representative firm may lead to unexpected macro outcomes owing to selection and reallocation effects between firms.

The workhorse of productivity research is the Solow growth model, and its extensions that include endogenous growth features such as R&D and human capital (Solow 1958; Romer, 1989; Griliches 1957). The growth model assumes profit maximizing behavior and perfect competition in the market for factor inputs, and is thus well grounded in micro-theory. The basic specification of the production function (in logs) is given below. Output is a function of the 'knowledge stock' or state of technology, A, and traditional factor inputs (capital, labor, energy, materials, services) weighted by their output elasticities which generally are restricted to add to unity.

$$y_i = A_i + \sum_{s \in klems} \alpha_s X_s$$

In the endogenous growth specifications, a second equation describes the evolution of the knowledge stock available for production as a function of explicit investments in knowledge or other intangibles and possible spillovers. Different versions of this model differ in de degree of spillovers ('standing on giants shoulders' vs 'fishing-out') and the effectiveness of own investment in increasing knowledge (see, e.g. Romer (1989), Segerstrom (1996), Jones and Williams (1998)).

$$\Delta A_i = G(I_i, A_i, \bar{A}),$$

where the additions to knowledge useable in production depend positively on own investment in knowledge, I, (e.g. R&D spending), the stock of locally available knowledge, A_i and spillovers from the stock of 'global knowledge', $\bar{A}$.

Using various versions of the underlying growth model, empirical work into the determinants of productivity proceeded along a few lines. One line of work used country/time panels (Barro; Mankiw, Romer, Weill) to look at issues such as convergence in productivity, the role of savings and the role of human capital. Related, researchers used industry/time panels for a single country to look at the effects of various factors other than capital and labor on productivity, for example R&D or infrastructure (Mohnen?, Aschauer, Fernald). Finally, the growth model is used as the theoretical background for measuring and quantifying productivity contributions to output growth, as done in growth accounting exercises.

To start with the latter, recently the EU 6th Framework program, EUKLEMS, delivered a cross-country, dataset on timeseries of output, inputs and prices by industry. The dataset also includes measures of the evolution of labor productivity and total factor productivity (TFP). The productivity numbers are derived by assuming constant returns to scale production technology and the equality of factor prices and marginal products of those factors in production. TFP growth then is the difference between the growth rate of output and a weighted sum of input growth rates, where the weights reflect the shares of revenue going as payment to each input. Solow notoriously called such a productivity index 'a measure of our ignorance', namely it is that part of growth that cannot be ascribed to the usage of inputs explicitly paid for by firms. The decisions of firms only pertain to the level of inputs and depend only on relative factor prices. The scale of operation is indeterminant, and productivity is exogenous to a firm's decisions.

Table 1 shows as an example the contributions to value added growth of ICT capital services and TFP for various aggregate sectors for the EU and the US in the previous decade. The table shows that output growth in the market sector in the US, from 1995 through 2005 was about 1.5 percentage point higher per year than in the EU. The lag in the EU not only occurred in the ICT producing sector, where the EU has relatively little activity in the very high growth industries, but was broad based in sectors where ICT adoption in the U.S. has increased significantly, such as the financial and distribution sectors. The contribution to output growth from the services provided by ICT capital was lower in the EU (0.4 percentage point) than in the US (0.6 percentage point), owing to a lower share of expenditures on ICT capital and a slightly lower rate of growth in ICT capital stocks. Half of the difference in output growth, however, is from lower growth of TFP, or that portion of output growth not explained by decisions of firms about the use of productive inputs.

The results from the growth accounting leave three unanswered questions. First, why is the contribution from ICT capital higher in the U.S.?, Next, why has TFP growth in the U.S. been so high, and what is the source of differences in TFP growth between the U.S. and the EU? Related, why is the share of the EU economy in the high growth sectors lower than in the U.S? The first question may be rephrased as "why is investment in ICT capital lower in the EU?" This is difficult to answer based on the available data on relative factor prices. In the next section we will extend the model to address this question. For the second question, the growth accounting method with its theoretical underpinning of profit maximizing representative firms in competitive factor markets, leads one to search for various types of 'externalities.' The third question falls outside of the scope of the model: In the growth model only the most productive ‘representative firms’ produces for the full market, or alternatively all firms are equally productive thus market shares become irrelevant.

Table 1.

1995-2005	EU			US		
	VA%	--Kict	--TFP	VA%	--Kict	--TFP
Market	2.1	**.4**	**1.0**	3.7	**.6**	**1.7**
EleCom	3.8	.8	2.8	10.5	.8	8.7
MfgxElc	1.2	.2	1.7	1.8	.2	2.2
DISTR	2.6	.3	1.5	4.1	.5	2.1
FinBus	3.5	.9	-1.0	4.3	.7	.4

Source: O’Mahony, Timmer, van Ark (2008).

The empirical search for externalities, or factors outside of traditional capital and labor inputs as a source of output growth, initially took place in a cross-country setting. In this empirical work many variables have been shown to have a significant correlation with output per capita, ranging from trade-openness, rule-of-law, colonial background, distance from the equator, among others. The end of the line for this type of inquiry must have been Sala-i-Martin, who explored all the possible combinations of potential externalities and shows results for those that jointly are correlated with productivity. While this provides some confidence on the selection of significantly correlated variables, it has provided no information for policy makers on how to boost, for example, growth in Africa, or on learning about any causal mechanisms that run from the explanatory variable to the resulting productivity increase.

Slightly more policy relevance was gleaned from single country industry/time panel datasets. For example, in the U.S., econometric analysis has measured the effect of R&D on industry productivity (for a review, see Griliches, 1998). In both the Solow-style and Romer-style growth models, R&D

expenditures build a knowledge stock that provides spillovers, thus they are a source of externalities and could potentially explain TFP growth. In general, there is empirical support for the idea that R&D boosts productivity and that the social rate of return to R&D is higher than the private rate, thus pointing to externalities and warranting some form of government action (see Jones and Williams 1998).

A more difficult task is to measure the effect of some policy related variable that varies over time, but not across industries. This would include most macro or structural policy variables, but also some externality-related input such as infrastructure. In a clever method of identification, Fernald (1999) shows that the highway system in the US is productive by establishing that it boosts productivity more in transport-intensive industries. Also shown, however, is that the effect of adding more highways is insignificant. Variants of this identification strategy have worked in establishing productivity effects of policy in cross-country settings, and will also be shown to be a useful method in utilizing micro-data to aid in empirical work in country/industry/time datasets.

In recent years, firm-level (panel) datasets have become available that allow researchers to assess how certain features at the firm-level, or various policy variables, affect firm-level productivity. For a variety of applications, see for trade (Bernard et al. 2007; Melitz 2003), ICT (Brynjolffson 2008; ONS 2008), FDI (Sabirianova et al. 2005), 'intangibles' investment (Bloom et al. 2007; Lynch 2007).

The advantage of the firm-level panels over more aggregated panel datasets, is that one actually observes decisions made by firms---the proper theoretical unit of observation. So, for example, the work by Bloom and Van Reenen (2006) show how certain types of management strategies affect the outcomes of a particular firm, whereas it is unclear that the effect would have been evident in a study using sectoral productivity and an average sectoral measure of management strategy. In recent literature, these productivity-enhancing strategies are considered to be investments in intangible capital that have not been accumulated into a measured capital input. If they had been capitalized, they would contribute to output growth in a manner similar to investment in tangible capital or in effective labor input. The results of this study, and similar studies showing the positive effect on productivity of firm-level use of ICT use (ONS 2008), do leave the question of why only certain firms seem to make these investments in intangibles. We return to this in the next section.

In the above examples of empirical work, the research attempted to show either how some variable was correlated with the TFP residual, or how some measure was related to quality of inputs. Examples of the latter are empirical estimates of the effect on output of increases in human capital, or growth in ICT capital services. Overall, this literature only allows two paths to higher growth: either through boosting the quality of productive inputs, or by providing some externality to the representative, or average firm. While many variables have been shown to be correlated with TFP, the empirical methods are not able to provide a prioritization nor is it easy to ascertain whether a policy change to provide more of the ‘externality’ actually would increase productivity at the margin. Further, the representative agent framework leaves out mechanisms for a wide variety of policies to potentially affect productivity.

3. Heterogeneous firms and productivity

The empirical productivity studies searching for sources of productivity growth using firm-level data are an improvement over sectoral and macro studies because the unit of observation, the firm, actually does face optimizing decisions that for the basis of the growth models. However, without further extension, the model does not take into account inherent differences across firms and the interactions between firms in a market. The underlying growth model is not able to cope with the observation that firms differ in size and in productivity, and that firms enter, exit, and grow at different rates. The following schematic the basic features of an extended growth model.

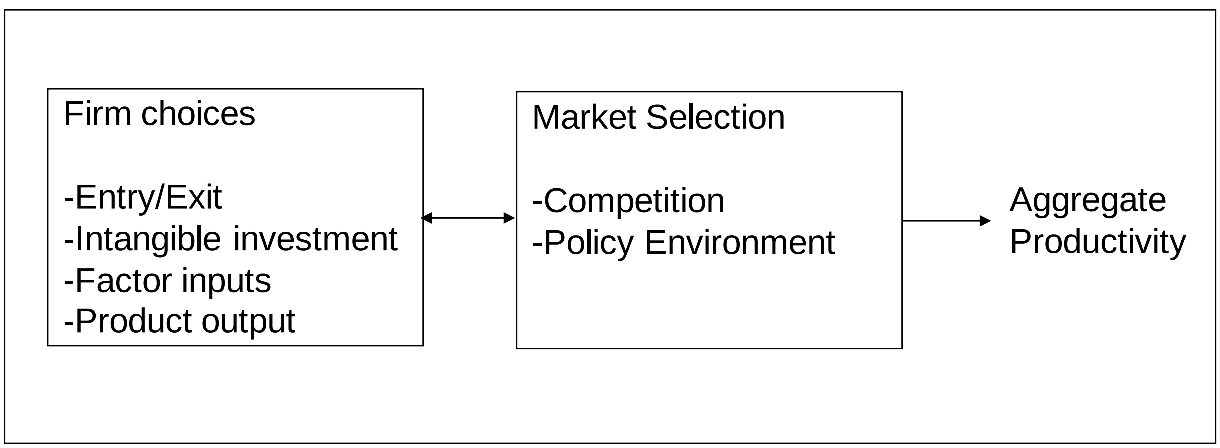

We start by showing some evidence of firm heterogeneity from the literature.. For this section, I liberally source from the previous literature, mostly papers by Bartelsman and various co-authors. The following section paraphrases or quotes directly from Bartelsman, Perotti, and Scarpetta (2008).

The stylized facts presented in this section are drawn from recent firm-level studies and from a harmonized database of indicators built up from firm-level data for a sample of OECD countries over the past decade.[2] The indicators for the OECD countries are generally limited to manufacturing industries and cover periods that vary by country but generally contain most of the 1990s. The reported facts relate to the size distribution of firms, the magnitude of firm entry and exit, the survival and post-entry growth of firms, and to the dispersion of productivity of entrants and incumbents. We start with a selection of indicators from the literature that portray the heterogeneity in firm characteristics and the amount of churn in employment and the population of firms.

-Over the first-half of the 1990s, firm turnover rates (entry plus exit rates) in OECD countries were in the range of 15 to more than 20 per cent in the business sector: i.e. a fifth of firms is either recent entrants, or will close down within a year.

-The process of entry and exit of firms involves a proportionally low number of workers: i.e. only about 10 per cent of employment is involved in firm turnover.

-Market selection is harsh in all countries. Only about 60-70 per cent of entering firms survive the first two years in the countries reviewed.

-Failure rates in the early years of activity are highly skewed towards small units, while surviving firms are not only larger, but also tend to grow rapidly.

-In the U.S. successful new firms expand rapidly compared with the EU. Bartelsman, Scarpetta and Schivardi (2004) show that the average size of surviving firms increases rapidly to approach that of incumbents in the market in which they operate.

-There is larger variation in the productivity levels of new firms in the U.S. than in Europe. The coefficient of variation (standard deviation divided by the mean) of the distribution of productivity levels of entrants varies across countries and manufacturing sectors. Results are reported in Bartelsman and Scarpetta (2004).

Table 2 shows stylized indicators derived from firm-level data, for the EU and the US, split by 'technology group'. The table shows indicators for total manufacturing, ICT-producing industries, and non-ICT industries. Omitted is information for ICT-using industries, whose indicator values always lie between the ICT-producing and non-ICT industries.

[2] The firm-level database collected indicators for 24 countries (Canada, Denmark, Germany, Finland, France, Italy, the Netherlands, Portugal, United Kingdom and United States Estonia, Hungary, Latvia, Romania, Slovenia; Argentina, Brazil, Chile, Colombia, Mexico, Venezuela, Indonesia, South Korea and Taiwan (Turkey, China coming)). These indicators are based on a process that involved the harmonization of key concepts (e.g. entry, exit, or the definition of the unit of measurement) as well as the definition of common methodologies for studying firm-level data. The methodology for collecting the country/industry/time panel dataset built up from underlying micro-level datasets has been referred to as 'distributed micro-data analysis' (Bartelsman 2004). A detailed technical description of the dataset may be found in Bartelsman, Haltiwanger and Scarpetta (2004).

Job destruction rates in the OECD countries hover around 10 percent.[3] The first row of Table 2 shows the percentage of job destruction occurring through the exit of firms rather than shedding of workers at continuing firms. Overall, this share is lower in the US than in the EU.[4] Most striking is the sizable difference in this rate in moving from the low to the high technology group in the U.S. In the high group in the US less than 7 percent of job losses occur through firm exit, while in the EU a third of the losses occur through firm exit. As a consequence, the high technology group releases a large quantity of employment at firms that continue to search for a fit in the market. These resources may be precisely those that are scarce at the more successful experimenting firms. By contrast, in the EU these resources remain attached to the firm, until the firm finally exits the market altogether. A similar pattern of job creation at entering firms versus total job creation emerges between the US and the EU and across technology groupings.

Table 2. Firm-level indicators by ICT-Technology Group

	US	EU	US	EU	US	EU
(percent)	Average Manuf		ICT Producing		Non-ICT	
Exit share of Job Destr.	24.7	34.3	10.7	24.1	24.9	37.4
Entrant Size rel. to incumbent	21.0	38.6	6.3	35.7	24.0	40.8
Productivity Gap of Exiters	10.0	15.4	1.2	9.1	7.9	17.7
Employment Share of Exiters*	18.9	23.1	20.2	31.8	19.8	22.3
Employment growth, top qrt.	68.6	50.1	91.8	65.1	70.8	45.0

The ICT-using industry is omitted from table. *The employment share of exiters is for 5-year window.

Entering and Exiting firms in the US tend to be smaller relative to industry average than in the EU. The relatively small size of entering and exiting firms in the US, especially in the high technology grouping, point towards the relative ease with which the mostly young firms can adjust their workforce to market circumstances. Further, as described in Bartelsman et al. (2005) and Aghion, Fally, and Scarpetta (2007), financial conditions and social safety nets may be such that firms with less certain 'business models' enter smaller in the US but then face a more vigorous selection process over time.

[3] We follow the Davis and Haltiwanger (1999) definitions of jobs creation, destruction, and gross flows, as described in BHS (2004).
[4] The EU countries used to compute these moments vary across indicators. The overall dataset includes information for Denmark, France, Finland, Germany, Italy, the Netherlands, Portugal, Sweden, and the United Kingdom.

The productivity threshold for exit is lower in the EU than in the US. The gap in productivity between exiting firms and incumbents is fifty percent larger in the EU than in the US. The ‘shadow of death’ (Griliches and Regev, 1995) of firms documented in the US is likely more pronounced in the US than in the EU. As firms face difficulty matching their production process to market demand, they shrink. This shrinking, all else equal, raises their measured productivity, but maybe not enough to remain competitive in the market. Conditional on exit, firms are seen to have a recent history of downsizing, but downsizing firms do have a higher chance of regaining productivity and market share.

Fewer resources are held in exiting firms in the US than in the EU. The share of employment taken up by firms that exit is lower in the US than in the EU, and lower for high technology firms. The measure presented here is the employment in year t-5 of firms that exit between year t-5 and t, as a share of total employment in year t-5.

Fast growing firms grow faster (and shrinking firms shrink faster) in the US. Over five-year periods, the average employment growth of high-tech firms in the quartile with the highest growth in employment in the US was 90 percent. This is higher than in the other technology groupings and higher than in the EU. The employment declines for the firms in the bottom quartile by employment growth over five years are a near perfect mirror image of the increases in the top quartile, both by country and technology groups.

Extending the Model

The above indicators clearly point to a large amount of variation across firms in productivity, size, and decisions to enter and exit. Starting with Baily, Hulten Campbell (1992), but also work of Bartelsman and Dhrymes (1998), and Foster, Haltiwanger and Krizan (FHK, 2001), researchers have prepared methods to decompose productivity growth into contributions from entry, exit, reallocation of resources between incumbents, and within-firm productivity growth. A commonly used decomposition, FHK, has shown its usefulness during periods of economic transition (Brown et al. 2006), but have otherwise been difficult to interpret. While simple reflection may make one think that a large contribution from exit is an indication that markets are functioning ‘properly’, either a high or a low exit contribution could coincide with higher economic welfare. Is it better for firms with declining productivity to remain large and in operation until the productivity gap with incumbents becomes very large---leading to a large productivity contribution---, or is it better for low productivity firms to shrink smoothly and exit rapidly, with a small productivity contribution ? Without a model, including adjustment frictions it is hard to make a welfare comparison.

An extension to the growth model is needed to take into account the heterogeneity and firm dynamics. To this end, we use the decomposition provided by Olley and Pakes (1998), which is a static decomposition of productivity levels. At the firm-level, the earlier presented model continues to hold. However firms first must decide upon their operating status, S, namely whether to enter (N), exit (E), or produce (C, continue). Next, they must make decisions regarding their factor inputs, including scale of operations, and their investment in knowledge-enhancing activities. In general, the production function is not assumed to be constant returns to scale. All these decisions depend not only on factor and output prices and the state of 'externalities', but also on (a firm's expectations about) decisions made by other firms.

$S_i \in \{N, E, C\}$

conditional on C:

$y_i = A_i + \sum_{s \in klems} \alpha_s X_s$, where $i \in C$

$\Delta A_i = G(I_i, A_i, \bar{A})$

and aggregate productivity

$\bar{A} = \sum_{i \in C} A_i + \sum_{i \in C} (\phi_i - \bar{\phi})(A_i - \bar{A})$

The final equation shows that aggregate productivity is a sum of average productivity of producing firms, A_i, plus the so-called Olley-Pakes cross-term. The latter term measures the covariance between firm size and firm productivity. Aggregate productivity is boosted if an above-average productivity firm is above average in size. Note that between periods, aggregate productivity changes not only because of a change in the two terms shown above, but also because the set of operating firms differs between periods owing to entry and exit.[5]

This simple framework allows many theoretical pathways by which policy and economic environment may affect productivity. In addition to policy that could alter factor prices (e.g. increasing the availability of high-skilled labor), boost innovative activity (promoting R&D), or provide external effects (public R&D, infrastructure), now policies may affect the 'status' decision, the allocation of resources across operating firms, as well as incentives for innovative investment.

The next sections show examples of how to build models and use empirical evidence to understand the various pathways of policy to aggregate productivity.

[5] For ease of exposition, we ignore the detail of timing usually present in these models. Assume that at the start of period t, new firms decide whether to enter, and then these entrants plus continuers from period t-1 decide whether to exit. Those that do not exit are continuers in period t. So in our notation, firms deciding on status E, must subsequently decide between X and C.

4. Examples: heterogeneous firms, policy, and productivity

Using the above framework, we will show examples of the different paths through which policy or economic environment can affect productivity. First, we show an example of modelling and empirical evidence that point towards a relationship between the economic environment and aggregate productivity through the path of resource allocation across firms. This example also displays how the selection margin is affected by policy, and in turn affects aggregate productivity. Next, we show examples of policy affecting selection of firms that differ in their innovative behaviour. This in turn, not only affects average firm-level productivity, but also aggregate productivity through the path of resource allocation. Finally, the paper presents an empirical example of how to use the underlying theory and firm-level data to show the policy effect on firm-level decisions to adopt ICT, through the paths of selection and resource allocation.

Reallocation and Productivity

In Bartelsman, Halitwanger, and Scarpetta (2008), our first example, we consider policy distortions to optimal resource allocation, as measured by the Olley-Pakes cross-term (OP-gap). In figure 1, a measure of the OP-gap is shown for manufacturing in a selection of countries in the 1990s. Among the OECD countries, aggregate productivity is nearly fifty percent higher than average in the U.S., while in mainland EU countries the OP-gap shows a twenty percent difference.

Interestingly, the new EU economies show a very low OP-gap. If resources had been allocated across firms with differing productivity with the role of the die, the measure would be zero. In figure 2, we see the OP-gap for a selection of transition economies over time. The overall pattern is that of large improvements over time in the measure of resource allocation.

Allocative efficiency (Olley Pakes decomposition -- cross term)
(weighted averages of industry level cross terms from OP decomposition)

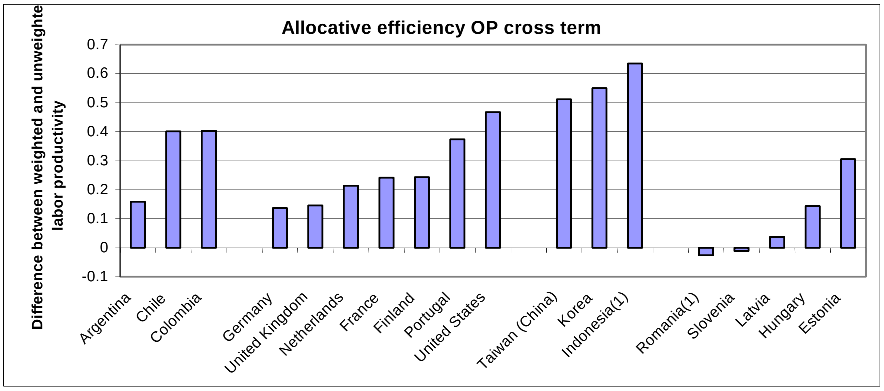

1. Based on the three-year differences

48

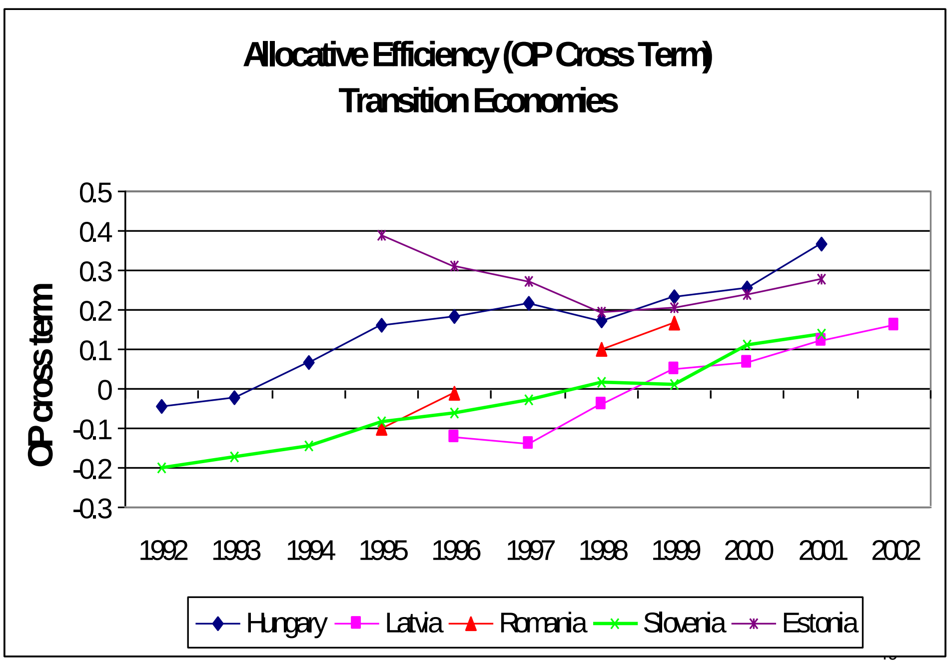

In Bartelsman, Haltiwanger, and Scarpetta, a model is built and simulated that attempts to explain the time-series pattern of the OP-gap in transition economies resulting from a shift in one policy parameter, namely one related to 'idiosyncratic distortions' to profitability of firms. An example of such a distortion would be business taxation that is not enforced uniformly across firms. Another possibility would be varying access across firms to 'external' factors, such as infrastructure or a key resource. In the model, the higher the variance in the distortion across firms, the lower is the OP-cross term.

Key elements of the model are that firms must pay an entry fee in order to learn about their productivity and their idiosyncratic tax or distortion (both drawn from random distributions). Firms then decide whether to exit or produce at their optimal scale.[6] On balance, the market is in equilibrium when the expected net present value of operating profits of continuing firms, corrected for exit, covers the entry fees. Without distortions to profitability, the optimal scale is monotonically increasing with the productivity draw, thus leading to a positive OP-gap.

Once an idiosyncratic tax is introduced, the tight link between productivity and size is disturbed. Productive firms with a bad tax draw remain inefficiently small, and firms that receive a subsidy become inefficiently large, thus reducing the OP-gap. The effects on aggregate productivity are less straightforward. The reduction in the OP-gap lowers aggregate productivity. However, selection plays a role as well. Because on average, firms facing high taxes are more likely to not produce, there will be a net subsidy to producing firms. This induces excessive entry, but also higher exit and thus more costly 'churn'.

In simulations where the tax is positively correlated with productivity, for example because the tax authority has been instructed to spend more resources to audit productive firms, the selection margin is particularly harmful to both resource allocation and to the characteristics of producing firms. On balance, highly productive firms will be less likely to produce and those that do will be inefficiently small. In future research, the type of model used in BHS also could be used to explore the effects of differing entry costs on aggregate productivity and welfare.

A key message from this work however, is that information on the evolution of firm-level productivity, the distribution of resources and market shares across those firms, and the process of entry and exit provide much information to analyse the effects of changes in economic policy. Just looking at aggregate outcomes is not enough. The flip-side is that on their own, indicators of entry and exit, or of the developments of resource allocation are not enough to make welfare comparisons because the entry, exit, and reallocation needed to improve efficiency also entail costs that are not fully policy induced.

[6] In order to support an equilibrium with firms differing in productivity and in scale, we assume a production function with decreasing returns to scale. Given the entry fee, this gives each firm its own U-shaped average cost curve and optimal scale. This specification borrows from Hopenhayn and Rogerson (1992). Alternatively, the demand side could have 'curvature' such that firms with differing productivity could be supported in an equilibrium (see e.g. Hsieh and Klenow (2007)).

Exit costs and innovation

A next example of a model with heterogeneous firms affected by policy is provided in a paper by Bartelsman, Perotti, and Scarpetta (2008). In their model of 'experimentation', firms must choose between risky but highly productive innovation, or safe innovation. The model, however, departs from the standard assumptions underlying the specification of knowledge creation in endogenous growth models. There, the uncertainty underlying innovation may be assumed away owing to risk pooling among investors, or may be added as a risk premium to the investment cost. In the experimentation model, firms must attach productive resources to the innovative activity. These resources need to be reallocated in case the developed 'technology' does not appeal to the market. As such, firms face costs in case their experiment is unsuccessful. In this model, it turns out that the incentive to participate in risky innovation is reduced by high exit costs. Further, aggregate productivity also is reduced by exit costs.

Firms may enter an innovation game by spending a fixed fee, I. Firms may choose 'safe' innovation, where their productivity, or profitability, is given by π, or they may choose a risky sector where their outcome is Π, with probability p. In the model, it is assumed that higher risk, or lower p, leads to higher payoffs. If firms attempt a risky innovation but fail, they pay a partial exit fee P_x and may try to innovate again, with a reconfiguration of resources. If they succeed, again with probability p, they obtain Π, otherwise they pay a total exit fee T_x. Figure 3 shows a schematic of the decision problem, where firms decide between receiving π or the expected value of the risky choice, which is $p\,\Pi + (1 - p)(-P_X + p\,\Pi) + (1-p)[-(1-p)P_X - T_X]$. In this model, a threshold level of exit costs can be computed such that no risky innovation will take place at higher exit costs. Further, the lower the probability of success, or the higher the benefit of success, the lower the exit cost threshold that chokes off innovation.[7]

This highly stylized model shows how firm-level choices depend upon macro policy, in this case exit costs. The heterogeneity across firms results from the idiosyncratic technological and market risk faced by firms. In some sectors or technology/market areas the difference between standard operations, π, and experimentation, Π, is small and the chance of succeeding, p, high, while in other areas the potential benefit of experimentation is quite high.

Identification of the hypothesized effect of exit costs, namely that it chokes off experimentation especially in the risky areas, using a single country firm-level dataset is problematic. All firms in the economy faces the same policy-induced exit costs, so no control group exists. Splitting firms across sectors based on ex-ante riskiness could help in identification if one could assume that the relationship between π and Π did

[7] For values of $p<0.5$.

Figure 3.

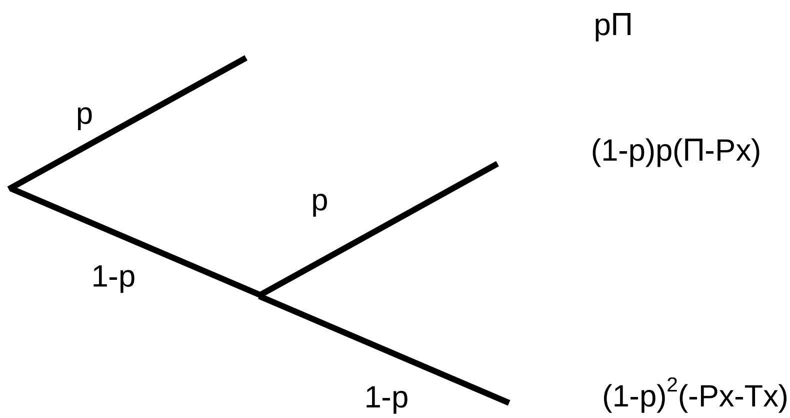

not vary across sectors. Time series variation in exit costs could help, but only if no other things change in the macro environment or if some valid instruments are found. This is the basic conundrum facing any researchers trying to identify the effects of some macro policy change using micro data: the thousands of observations don't really help the basic fact that the identification depends on two periods: before and after the policy change.

Instead, using country/time/industry datasets, the cross-country variation in exit costs along with controls for fixed effects provide the means to identify how much exit costs harm productivity in low 'p' relative to high 'p' industries. In the Bartelsman et al. (2008) paper, empirical exploration of the hypothesized effect of exit costs is conducted using the EUKLEMS dataset, which is an industry/country dataset of output and input indicators. Further, various indicators of exit costs, provided by the OECD or the World Bank, are used that vary by country (and over time, albeit not very reliably measured). These indicators include measures of stringency of employment protection legislation (EPL, from OECD, used e.g. by Bassanini et al. 2008) or costs of exiting (from World Bank 'Cost of doing business database', used e.g. by Djankov et al. (2002)).

The remaining piece of information requiremed for empirical testing is an indicator of riskiness or potential benefit of innovation in an industry. In their empirical work, Bartelsman, Perotti and Scarpetta use various indicators to rank the inherent riskiness of innovation in a sector. Because observables used to indicate riskiness themselves are affected by exit costs in a country, the indicators are drawn from the U.S. or the U.K, which are both countries with low exit costs. One indicator of industry rank is derived from a measure of how much higher productivity in the best quartile of firms in an industry is relative to the average. Another indicator is related to the penetration of new technology in an industry, namely the diffusion of broadband internet among firms in an industry for the U.K.

The implication of the simple model is that in sectors where the probability of success is low, or the gain from experimentation is high, exit costs will reduce aggregate productivity more than in sectors where the gains from experimentation are lower. Table 3 shows some selected regressions from Bartelsman et al. (2008), using the following regression equation:

$$v_{ijt} = \sum_{s \in k^N k^{IT} l} \alpha_s X_{s,ijt} + \gamma_c I_{jt} + \gamma_s I_{jt} F_i + FE + \varepsilon_{ijt}$$

Value added in an industry (i), country (j) and time (t) is regressed on traditional inputs (ICT-capital services, non-ICT capital services, and labor hours), on the measure of exit costs (I), sometimes interacted with the industry rank in the riskiness measure (F), and on different specifications of fixed effects (FE). In the first column of table 3, in a regression using the index of employment protection legistlation (EPL) as the measure of exit costs and no interaction, we see that on average EPL does not significantly affect productivity. In the next columns, the country and time varying EPL indicator is interacted with the industry rank of riskiness. In all the cases, EPL significantly reduces productivity of high risk sectors relative to low risk sectors. In the Bartelsman et al. paper, the coeficient of exit costs interacted with industry rank is significant for many combinations of exit cost measures, industry ranking measures, subsets of countries, and differing time periods.

Table 3

Dependent var: / Regressor:	Log(VA)	Log(VA)	Log(VA)
Log: Kit,Knit,Hours	***	***	***
EPL	.47	.34	.46
	(0.02)	(0.14)	(0.19)
EPL x Rank	---	-1.18	-1.13
		(3.07)	(3.08)
Rank variable	---	Top quartile prod/mean	Broadband-use
Num. obs.	7032	6790	7031
R-sq	.97	.97	.97

Source: EUKLEMS, BPS(2008); 1991-2004; Fixed effect: country, industry, time. T-stats in parenthesis.

Exit costs and employment

Above, we saw that productivity levels are lower owing to high exit costs precisely in those industries where the most could be gained through experimentation with new technologies or new ways of meeting market demand. The model only describes the endogenous selection of firms that differ in their innovation strategy, and is not rich enough to make predictions about the effect of exit costs on resource allocations across sectors, or between firms in an industry conditional on their innovative behavior. In preliminary explorations, Bartelsman, Gautier, and de Wind (2008, paper not yet available) are building an equilibrium labor search model of the Mortenson-Pissarides type (1994) where vacancies may be opened in a safe or a risky sector. A riskier sector is one where the fulfilled job receives a productivity draw from a wider distribution than in a safer sector, although the mean draw may be the same. A job match that gets terminated must pay a policy-determined exit cost. In this model, the overall sign of the effect of exit costs on total employment and unemployment depends on the elasticity of moving out of home production into the labor force. However, the labor share of employment in the riskier sectors is harmed by exit costs, while the labor share going to safer sectors increases. The riskier the sector, the larger the negative effect of exit costs on labor share.

Using the EUKLEMS dataset, the share of employment in a country allocated to a particular sector is regressed on an indicator of exit costs, the same indicator interacted with the rank of industry riskiness, and various fixed effects. One set of results is shown in table 4 below.

Table4

Dependent var: Regressor:	Labor share in sector	Labor share in sector
EPL	.02	.02
	(0.74)	(0.74)
EPL x Rank	-0.82	-0.84
	(10.30)	(10.55)
Rank variable	Top quartile prod/mean	Broadband-use
Num. obs.	5518	5518
R-sq	.84	.84

Source: EUKLEMS, BPS(2008); 1995-2005; Fixed effect: country, industry, time. T-stats in parenthesis.

As seen, the effect of EPL on employment share in an industry is positive for safe industries, and negative for risky industries. The results hold up for a wide range of country sub-samples, time periods, exit cost indicators and rankings of riskiness of industry.

[This paragraph needs updating….. re-do numerical calculations, using same sample period/countries/industries in both tables] With estimates of the effect of exit costs on employment and on productivity by industry, an estimate can now be made of the effect of a change in exit costs on aggregate productivity. However, this is a 'ceteris paribus' experiment, and does not take into account possible interactions between employment and productivity changes jointly. Nonetheless, in an industry with a high 'riskiness', say retail trade, a reduction in the employment protection index from 4 to 3, (say from Germany to the Netherlands) may increase productivity by one percent and increase employment share by xx percentage points. Integrating over all industries, overall productivity would increase by xx percent.

$\Delta\Pi = \sum_i \Delta\pi_i\phi_i + \Delta\phi_i\pi_i$, assuming $\sum_i \Delta\phi_i\Delta\pi_i = 0$.

It is important to note, however, that the increase in aggregate productivity that may be obtained through reducing exit costs does not necessarily mean an increase in welfare. In both models discussed above, an increase in resources devoted to the risky sectors necessarily means a higher rate of 'churn' or reallocation of resources. To the extent that these entail inherent costs of such as search costs, sunk investments, etc, the improvement in productivity comes at a price. It is beyond the scope of this paper to provide estimates of these costs.

Experimentation and Intangible Investment

In the first section of this paper, it was seen that the representative agent models were unable to provide a satisfactory answer as to why the EU had lower investment in ICT and fewer resources devoted to high growth ICT sectors. Further, empirical findings of the productivity effects of intangible investments, in ICT but also in organization and management, showed a positive effect on productivity at the firm level. However, the empirical work was not able to uncover the drivers of intangible investment or of policies that may affect these investments. The previous sections, by contrast, showed the incentives that would induce or inhibit firms to allocate resources into risky sectors. In both models, a fixed investment would buy an option on a risky outcome. Conditional on success, the resources devoted to the operation would increase, thereby leveraging the original investment, while conditional on failure, the resources would be redeployed elsewhere. Exit costs would make such an option more costly.

Investments in intangibles, such as software, new business methods, new means of meeting customer demands, etc, fit into this type of risky investment. The investment must be implemented in a firm operating in the market. If the 'system' works, the firm can duplicate the new software or business method in other production units without full duplication of investment expenditures. Some costs, such as costs of training workers to use the new business methods, or costs of complementary hardware, are added with replication, but a certain portion remains 'non-rival' in production and thus leads to declining average costs as the firm scales-up. This story fits very well with the case-study by Brynjolfsson, McAfee, Zhu and Sorell (2006), of a firm designing, implementing, and rolling-out an ICT-based innovation in business process.

The Brynjolfsson et al. finding that ICT investment goes together with increases in variation in firm-growth is corroborated in a recent Eurostat project 'ICT Impact Assessment' (ONS, 2008). In chapter 13 of this report, written by Bartelsman, the following scatterplot shows the correlation between the percent of workers in an industry in a given year and country with access to broadband internet and the width of the distribution of firm-level value added growth growth rates. More 'churn' in industry market share is correlated with adoption of advanced, risky, technology, as seen in Figure 4.

Figure 4.

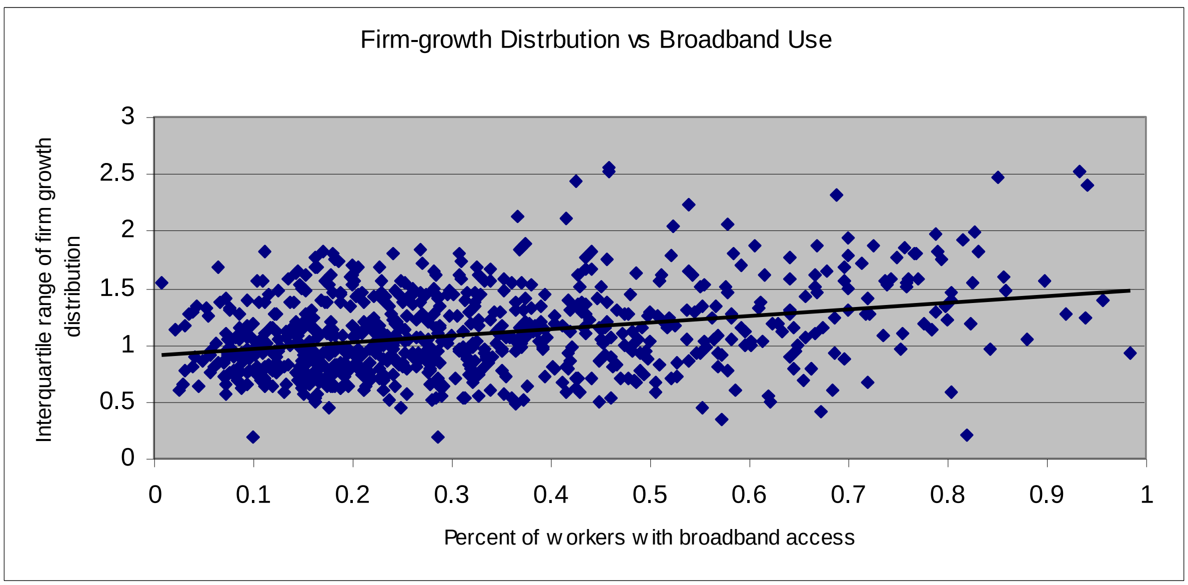

The Eurostat report further shows the results of an empirical exercise with a production function and an ICT adoption equation. Value added (in logs) in an industry/country/year depends on the percentage of workers with broadband access, traditional factor inputs and fixed effects.

$$a: v_{ijt} = a_0 + a_1 DSL\% + a_2 k^{IT} + a_3 k^N + a_4 hrs + dummies$$

$$b: DSL\%_{ijt} = b_0 + b_1 w_{-1} + b_2 Cap\%_{-1}^{IT} + b_3 HiSkl_{-1} + b_4 Churn + dummies$$

The production function etstimates alone show a positive impact of broadband use, over and above the contribution of the other factors, including ICT-capital services. Because the broadband use may be endogenous, and because policy makers want to know if a policy lever could increase broadband use at the margin, it is preferable to jointly estimate an adoption equation. Often, technology adoption is modelled by looking at costs and benefits. Another way this is described in the literature, is by looking at the ability of firms to adopt and their desire to adopt. Typically, the variables hypothesized to affect adoption of a technology would include its price, the ability and readiness of the firm to take on the new technology (in our case, skilled ICT workers, and availability of ICT capital), as well as some measure of expected benefits. Given the non-rival, but likely appropriable, nature of firm-level ICT projects, firms that can successfully replicate, or scale-up, their operations upon success in the market, have a higher potential benefit to the ICT investment (see e.g. Brynjolfsson et al 2007, Bartelsman et al. 2008). To proxy this, we use the width of the distribution of firm-level output growth.

Table 5

Coef	Variable	a:Log (value added); b:DSL%	
a1	Broadband Penetration (DSL%):	1.24	.90
a2	Non-ICT Capital	.35	.27
a3	ICT Capital	-.07	.05
a4	Labor Hours	.72	.68
b1	Wage(t-1)	.24	.02
b2	ICT capital share(t-1)	.31	.20
b3	High-skill labor share(t-1)	.18	.38
b4	Churn (interqrtl range of firm growth distribution)	.30	.15
	Fixed effects	c,t	i,t
	Num. Obs.	659	646

Source: EUKLEMS and ONS; 2001-2005; All coefficients are statistically significant at 5%-level.

The production and adoption equation are estimated jointly in a system that uses lagged exogenous variable as instruments. Table 5 shows coefficient estimates in versions of the two equations, with the (log) of value added as the dependent variable of the production equation (a), and the percent of broadband enabled workers as the dependent variable in the adoption equation (b). The system is run either with country and time fixed effects, or with industry and time fixed effects. As shown, the broadband access variable has a significant positive effect on TFP. The ICT capital variable becomes insignificant (and sometimes negative), while the hours and non-ICT capital variables are as expected. The explanatory variables in the adoption equations also show up significantly positive in all specifications. In the ONS (2008) report, an alternative specification also is shown, with the production equation run in (log) first-differences. Because estimation of capital output elasticities is particularly sensitive to the fact that capital is subject to steep adjustment costs, a combined 'input growth' variable is constructed from the separate traditional inputs, using expenditure shares as weights. With this specification, the percentage of broadband enabled workers is seen to boost output growth by one-tenth to one-quarter of a percentage point, while the explanatory variables in the adoption equation continue to be significantly positive.

5. Conclusions and directions for future research

A broad range of research over the past decades has attempted to provide empirical evidence on the drivers of productivity growth. The broad range of methods and data sources used in the research provides policy makers with a diffuse and ambiguous set of findings to be used for formulating policy priorities towards enhancing productivity growth. Until about a decade ago, the empirical evidence was based on macro or industry-level aggregate data. The growth models used in the research however were based on micro-economic decisions and applied to aggregate data by relying on the representative firm paradigm. The research using aggregate data has been able to provide evidence on the impact of R&D on productivity, but could not identify the role of policies that did not impact productivity either through spillovers or through improvements in factor input quality.

In the past decade, firm-level data has become available in many countries. These data allow researchers to avoid having to reply on the representative firm paradigm and have the correct unit of observation needed to observe economic behaviour. Further, the micro data allow one to extend the underlying model to include new routes through which policy can affect productivity, namely through selection and through resource (re)allocation. However, single-country firm-level data generally do not provide the information necessary to identify the impact of macro policy. The reason is that the policy affects all firms, leaving only the 'before/after' variation to identify the impact of a policy change.

Having access to cross-country firm-level data would allow identification of policy affects through variation across countries in the policy stance and in the timing of policy changes. Recently, using publicly available data of large firms (e.g. the Amadeus database), researchers are exploring this route.

However, the comprehensive surveys and censuses of enterprises, firms, and plants available at national statistical offices generally are confidential and cannot be used in a cross-country manner. In the examples shown in this paper, the identification occurs by using industry (aggregate) data that has been enhanced by adding indicators built up from confidential firm-level data available in each country. The identification occurs through cross-country variation, but also through the differential impact that policy may have based on observable (with micro data) characteristics of the industries. Using this method, the paper shows how exit costs affect productivity and employment in sectors that rely have the potential of large but risky gains through innovation.

In the future, similar empirical approaches could be used to identify the impact across EU countries of diverse policies ranging from changes in marginal tax rates, broader labour market policies, worker training, intellectual property rights, competition policy, infrastructure expenditures, etc. For each of the particular policies one would need to find, based on theoretical grounds, indicators from the micro data that rank the industries according to the potential impact of the policy. This work entails building models appropriate for the issue at hand that feature heterogeneous firms, market interactions and selection, as well as dynamics of allocation of resources and market shares across firms. Next, it requires tapping into firm-level datasets in a large set of countries to extract the required indicators. The last step, however, need not be very sophisticated, as the examples in this paper show.

References

Aghion, Philippe, Thibault Fally, and Stefano Scarpetta. 2007. Credit constraints as a barrier to the entry and post-entry growth of firms. *Economic Policy* 22, no. 52: 731-779. doi:10.1111/j.1468-0327.2007.00190.x.

Baily, Martin Neil, Charles Hulten, and David Campbell. 1992. Productivity Dynamics in Manufacturing Plants. *Brookings Papers on Economic Activity. Microeconomics* 1992: 187-249.

Barro, Robert J. 1991. Economic Growth in a Cross Section of Countries. *The Quarterly Journal of Economics* 106, no. 2 (May): 407-443.

Barro, Robert J., and Xavier Sala-i-Martin. 1992. Convergence. *Journal of Political Economy* 100, no. 2 (January 1): 223. doi:10.1086/261816.

Bartelsman, E. J., and P. J. Dhrymes. 1998. Productivity Dynamics: US Manufacturing Plants, 1972–1986. *Journal of Productivity Analysis* 9, no. 1: 5-34.

Bartelsman, Eric, Stefano Scarpetta, and Fabiano Schivardi. 2005. Comparative analysis of firm demographics and survival: evidence from micro-level sources in OECD countries. *Ind Corp Change* 14, no. 3 (June 1): 365-391. doi:10.1093/icc/dth057.

Bartelsman, Eric J., John Haltiwanger, and Stefano Scarpetta. 2004. *Microeconomic evidence of creative destruction in industrial and developing countries*.

Bartelsman, Eric J., John Haltiwanger, and Stefano Scarpetta. 2008. *Cross Country Differences in Productivity: The Role of Allocative Efficiency.*

Bartelsman, Eric J., Enrico Perotti, and Stefano Scarpetta. 2008. *Barriers to Exit, Experimentation and Comparative Advantage, mimeo.*

Bassanini, Andrea, L. Nunziate, and Danielle Venn. 2008. Job protection legislation and productivity growth in OECD countries. IZA DP no. 3555, July.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen. 2007. Americans Do I.T. Better: US Multinationals and the Productivity Miracle. *National Bureau of Economic Research Working Paper Series* No. 13085 (May). http://www.nber.org/papers/w13085.

Brown, J. David, John S. Earle, and Almos Telegdy. 2006. The Productivity Effects of Privatization: Longitudinal Estimates from Hungary, Romania, Russia, and Ukraine. *Journal of Political Economy* 114, no. 1 (February 1): 61-99. doi:10.1086/499547.

Brynjolfsson, Erik , Andrew Mcafee, Feng Zhu, and Michael Sorell. 2006. Scale Without Mass: Business Process Replication and Industry Dynamics. *SSRN eLibrary*. http://ssrn.com/paper=980568.

Djankov, S., and R. La Porta. 2002. The regulation of entry. *The Quarterly Journal of Economics* 117, no. 1: 1-37.

Fernald, John. 1999. Roads to Prosperity? Assessing the Link between Public Capital and Productivity. *American Economic Review* 89, no. 3: 619--638.

Foster, Lucia, John C. Haltiwanger, and C.J. Krizan. 2001. Aggregate Productivity Growth: Lessons from Microeconomic Evidence. In *New Developments in Productivity Analysis*. Chicago: University of Chicago Press.

Griffith, R., S. Redding, and J. Van Reenen. 2003. Mapping the Two Faces of R&D: Productivity Growth in a Panel of OECD Industries. *Institute for Fiscal Studies and University College London; Centre for*.

Griliches, Zvi. 1957. Hybrid Corn: An Exploration in the Economics of Technological Change. *Econometrica* 25, no. 4 (October): 501--522.

---. 1998. *R&D and Productivity*.

Griliches, Zvi, and Haim Regev. 1995. Firm productivity in Israeli Industry: 1979-1988. *Journal of Econometrics* 65, no. 1 (January): 175--203.

Jones, C. I., and J. C. Williams. 1998. Measuring the Social Return to R&D. *Technology* 1119.
Lynch, Lisa. 2007. The adoption and diffusion of organizational innovation: Evidence for the U.S. Economy. CES-WP-07-18, June.
Mankiw, N. Gregory, David Romer, and David N. Weil. 1992. A Contribution to the Empirics of Economic Growth. *The Quarterly Journal of Economics* 107, no. 2 (May): 407-437.
Melitz, Marc J. 2003. The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. *Econometrica* 71, no. 6: 1695-1725. doi:10.1111/1468-0262.00467.
Mortensen, Dale T., and Christopher A. Pissarides. 1994. Job Creation and Job Destruction in the Theory of Unemployment. *Review of economic studies* 61, no. 208: 397-415.
Nicoletti, Giuseppe, and Stefano Scarpetta. 2003. Regulation, productivity and growth: OECD evidence. *Economic Policy* 18, no. 36: 9-72. doi:10.1111/1468-0327.00102.
Olley, G.Steven, and Ariel Pakes. 1996. The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica* 64, no. 6: 1263-1297.
Rodrik, Dani. 1996. Understanding Economic Policy Reform. *Journal of Economic Literature* 34, no. 1 (March): 9-41.
Romer, Paul M. 1990. Endogenous Technological Change. *Journal of Political Economy* 98, no. s5 (January 1): S71. doi:10.1086/261725.
Sabirianova, Klara, Jan Svejnar, and Katherine Terrell. 2005. Distance to the Efficiency Frontier and Foreign Direct Investment Spillovers. *Journal of the European Economic Association* 3, no. 2-3 (April 1): 576-586. doi:10.1162/jeea.2005.3.2-3.576.
Sala-I-Martin, Xavier X. 1997. I Just Ran Two Million Regressions. *The American Economic Review* 87, no. 2 (May): 178-183.
Scarpetta, S., T. Tressel, and Organisation for Economic Co-operation and Development Economics Dept. 2002. *Productivity and Convergence in a Panel of OECD Industries: Do Regulations and Institutions Matter?* OECD.
Segerstrom, Paul S. 1998. Endogenous Growth without Scale Effects. *The American Economic Review* 88, no. 5 (December): 1290-1310.